preliminary and slightly incomplete; comments are very welcome

Market Design in Regulated Health Insurance Markets: Risk Adjustment vs. Subsidies*

Liran Einav, Amy Finkelstein, and Pietro Tebaldi†

July 2018

Abstract. Health insurance is increasingly provided through regulated competition, in which two key market design instruments are subsidies for consumers and risk adjustment for insurers. Although typically analyzed in isolation, we illustrate through a stylized model that subsidies offer two key advantages over risk adjustment in markets with adverse selection: they provide higher flexibility in tailoring insurance premiums to buyers with different willingness to pay and, under imperfect competition, they produce equilibria with lower markups and greater enrollment. We provide a quantitative assessment of these effects using demand and cost estimates from the 2014 California exchange regulated by the Affordable Care Act. There, we estimate that holding government spending constant, subsidies can provide a Pareto improving increase in enrollment of about 6 percentage points, or 13 percent, over risk adjustment alone.

JEL classification numbers: G22, G28, H51, I13

Keywords: Health Insurance, Insurance Exchanges, Market Design.

*Einav and Finkelstein gratefully acknowledge support from the Sloan Foundation. Tebaldi acknowledges support from the Becker Friedman Institute. We thank Ben Handel, Mike Whinston, and many seminar participants for helpful comments.

†Einav: Department of Economics, Stanford University, and NBER, leinav@stanford.edu; Finkelstein: Department of Economics, MIT, and NBER, afink@mit.edu; Tebaldi: Department of Economics, University of Chicago, ptebaldi@uchicago.edu.

1 Introduction

For modern governments, social insurance – particularly pensions and healthcare – represent a large and growing fraction of expenditures. Historically, these were provided through direct public provision of insurance. Increasingly, however,social insurance involves public regulation of privately-provided products, rather than direct government provision. In the context of health insurance in the United States, examples include the approximately one-third (and growing) share of Medicare provided by private firms through the Medicare Advantage program, the 2006 introduction of Medicare Part D for prescription drug coverage provided by regulated private insurers, and, most recently, the regulated private health insurance exchanges created by the Affordable Care Act (ACA). The trend toward regulated competition in health insurance is not limited to the United States; in countries such as the Netherlands, Switzerland, and Chile, large components of their universal coverage are now offered via insurance exchanges that are similar to those created by the ACA.

In all of these cases, the market sponsor (typically a government) sets the rules, and private firms compete within the rules to attract enrollees. Once the sponsor has defined the set of insurance products that can or must be offered, there typically remain two key market design decisions: premium subsidies for consumers and risk adjustment for insurers. Subsidies are typically viewed as the instrument by which premiums are made affordable; they are therefore often linked to the income of potential buyers. Risk adjustment systems, by contrast, are typically viewed as a way to compensate participating insurers for enrolling high-cost buyers in order to reduce concerns about adverse selection and risk skimming.

Perhaps as a result, policy discussions and academic analyses typically study subsidies and risk adjustment in isolation, as two separate and unrelated objects. For example, recent work has focused on the impacts of risk adjustment on cream-skimming in Medicare Advantage (Newhouse et al. 2012; Brown et al. 2014) and on the impact of premium subsidies on enrollment by low-income individuals in ACA or similar exchanges (Finkelstein, Hendren, and Shepard 2017; Frean, Gruber, and Sommers 2017).

In this paper, we make the simple observation that these two instruments – often set by the same entity – naturally interact through their impact on equilibrium allocation. In the stylized price theory framework of Einav, Finkelstein, and Cullen (2010), subsidies are instruments that shift out the demand curve, while risk adjustments are instruments that rotate (and potentially shift) the cost curve. Given that equilibrium is determined by the intersection of demand and cost, it seems natural to study these two market design features in tandem, and to ask how they may interact, and the extent to which they substitute or complement each other.

We begin with a standard, highly stylized setting of an insurance exchange, with horizontally differentiated insurance plans and adverse selection (willingness to pay that is increasing in costs). In contrast to the way subsidies and risk adjustment are typically used in practice, we take a more conceptual approach and allow risk adjustment and subsidies to be functions of the same set of observables. Our main theoretical result is that, for a given level of market sponsor expenditures, subsidies can achieve higher enrollment and higher consumer surplus than risk adjustment; moreover, they can do so while making each type of consumer (weakly) better off. We discuss how this stems from two distinct forces. First, because they flatten the cost curve, risk adjustment increases equilibrium markups, a point that has been noted previously in the literature (Starc 2014; Mahoney and Weyl 2017). We show that, as a result, a (uniform) subsidy can achieve higher coverage and higher consumer surplus at a given level of sponsor spending than risk adjustment can. We refer to this as the "markup effect." Second, non-uniform subsidies can additionally reduce inefficiencies arising from adverse selection by targeting different (observable) types of buyers with different premiums, to incentivize low-risk types to purchase insurance. We refer to this as the "premium targeting effect."

The theory provides qualitative results in a simplified horizontally differentiated setting. To explore the comparison between subsidies and risk adjustment quantitatively in a richer setting, we use the estimates of demand and costs in Tebaldi (2018) from the first year (2014) of the Affordable Care Act (ACA) health insurance exchange in California. The California exchange is among the three largest ACA marketplaces, with 1.3 million enrollees. Over 90% of them received federal subsidies, with annual public expenditures totaling approximately 3.5 billion dollars. Tebaldi (2018) used the estimates from this setting to explore optimal subsidy design. We expand the analysis here to consider comparisons between optimal risk adjustment and optimal subsidies.

Holding fixed market sponsor spending, we estimate the enrollment-maximizing equilibrium – and quantitatively compare enrollment, markups, consumer surplus, and producer surplus – under different market regimes: risk adjustment, uniform subsidies, and non-uniform subsidies. The enrollment increase between optimal risk adjustment and uniform subsidies reflects the markup effect; the enrollment increase between uniform and non-uniform subsidies reflects the premium targeting effect. Of course, such premium targeting may make higher-risk types worse off, creating a social trade off between maximizing enrollment and type-specific utility. We therefore also consider potential gains from optimal non-uniform subsidies that – relative to the optimal uniform subsidy or risk adjustment – make all types (weakly) better off. At roughly current levels of market sponsor expenditures, we estimate that optimal subsidies increase enrollment by about 6 percentage points (13 percent) compared to optimal risk adjustment. In our setting, this is primarily driven by the markup effect, with less than a third coming from the premium targeting effect. We also find that in our setting there is little or no penalty on enrollment from imposing the Pareto restriction on the optimal non-uniform subsidy.

Our paper is related to an increasing number of papers that study subsidy design (Chan and Gruber 2010; Decarolis, Polyakova, and Ryan 2017; Finkelstein, Hendren, and Shepard 2017; Jaffe and Shepard 2017; Tebaldi 2018), risk adjustment (Glazer and McGuire 2000; Van de Ven and Ellis 2000; Ellis 2008; Newhouse et al. 2012; Brown et al. 2014; Einav et al. 2016; Layton, McGuire and Sinaiko 2016; Geruso and Layton 2018) , the ACA exchanges (Dickstein et al. 2015; Abraham et al. 2017; Frean, Gruber, and Sommers 2017; Tebaldi 2018), and the design of health insurance exchanges more generally (Curto et al. 2014; Handel, Hendel, Whinston 2015; Azevedo and Gottlieb 2017). As noted earlier, in all these papers subsidies and risk adjustments are treated in isolation, and none of these papers engages in the interaction between them.

Our paper proceeds as follows. In Section 2 we present a stylized theoretical framework and results. Section 3 briefly describes the empirical setting and estimates, which are taken directly from Tebaldi (2018). Section 4 presents our quantitative assessment of alternative market regimes. In addition to our "bottom line" numbers for the entire California market, we also report some comparative statics in a more simplified setting in order to illustrate in a more transparent way some of the underlying economic forces at work. The last section concludes with a brief discussion of potential reasons why, despite our theoretical and empirical results, risk adjustment remains an increasingly popular market design instrument.

2 A stylized theoretical framework

2.1 Setting and notation

We consider a stylized setting. There is a single insurance coverage contract, offered by J competing insurers, each indexed by j. As is often the case in regulated insurance markets, insurers are not allowed to charge different prices to different consumers (beyond any price discrimination that is built into the subsidy design), so each insurer j sets a single price p_j in a Bertrand-Nash Equilibrium.

Potential buyers are heterogeneous, with each consumer i defined by a triplet (v_i, c_i, w_i). $v_i = (v_{i1}, ..., v_{iJ})$ is a vector of consumer i's willingness to pay for the insurance contracts offered by the different insurers. We denote by $c_i > 0$ the expected cost to the insurer of covering individual i, which for simplicity we assume to be the same across insurers. Finally, w_i denotes a vector of observable characteristics, such as age, income, or health risk score, which in principle can be used as an input to the subsidy or risk adjustment design. We also refer to w_i as consumer i's type, and take the number of types as finite throughout.

In this setting, similarly to the framework in Einav, Finkelstein, and Cullen (2010), the population is represented by the joint distribution of v_i, c_i, and w_i. This imposes no restrictions on the relationship between preferences, cost, and consumer types, but we note that it is natural to think of adverse selection as a positive correlation between cost c_i and willingness to pay v_{ij} (for each j). When this is the case, if $ac_j(p)$ is the average cost of individuals covered by contract j when premiums are $p = (p_1, ..., p_J)$, adverse selection implies that $\partial ac_j(p)/\partial p_j > 0$.

A subsidy design is defined as a function $s(w_i)$. If buyer i buys insurance coverage from insurer j, she pays $p_j - s(w_i)$, the market sponsor pays $s(w_i)$, and the insurer's (expected) profits from covering buyer i are $p_j - c_i$.

A risk adjustment design is defined as a function $r(w_i)$. If buyer i is insured, the market sponsor transfers $r(w_i)$ on top of the premium the insurer receives; insurer j's profits from covering individual i are therefore $p_j - (c_i - r(w_i))$. We note that we do not require risk adjustments to be budget neutral, so in principle risk adjustment payments could result in greater (or lower) overall expenditure by the market sponsor. Importantly, throughout the paper we restrict attention to only "ex ante" and "regular" risk adjustment designs, which are the most common in mature markets.

By "ex ante" we mean that the risk adjustment function associated with buyer i is known at the time of enrollment, and does not depend on buyer i's subsequent realized costs or realized costs of other buyers in the market.[1] By "regular" we mean that the risk adjustment reduces adverse selection by compensating insurers more generously for covering more risky buyers, or more precisely that $r(w_i) \geq r(w_k)$ if and only if $E[c_i|w_i] \geq E[c_k|w_k]$. Under a regular risk adjustment and adverse selection, the per-enrollee risk adjustment transfer received by j is increasing in p_j.

2.2 Perfect competition

We begin by analyzing the case of perfect competition, which arises in our setting when insurers are homogeneous ($v_{i1} = v_{i2} = \ldots = v_{iJ}$, for all i). As a result, as in Einav, Finkelstein, and Cullen (2010), the Bertrand-Nash Equilibrium implies that insurers set prices so that price equals average costs, and profits are zero.

In such a case, we obtain the following result.

Proposition 1 *Under perfect competition, for any Nash equilibrium that is achievable with risk adjustment, there exists a subsidy design with no risk adjustment that can achieve the same equilibrium, with the same enrollment for all types and the same total spending by the market sponsor.*

The proof is in the appendix. The intuition is simple and is illustrated in Figure 1. It plots demand and average cost curves as in Einav, Finkelstein, and Cullen (2010). The demand curve shows the share of the population who buy insurance at a given price. The average cost curve shows the average expected claims among all consumers who buy insurance at that price. Under perfect competition, the equilibrium is given by the intersection of the demand and average cost curve.

The left panel illustrates the impact of risk adjustment. Risk adjustment flattens – and potentially shifts – the average cost curve, leading to a new market equilibrium with greater insurance coverage and lower insurer prices. Proposition 1 implies that this market equilibrium can alternatively be achieved by reverting back to the original cost curve, and using a uniform subsidy to achieve an appropriate parallel shift of the demand curve; this is illustrated in the right panel of Figure 1.

Proposition 1 implies that, under perfect competition, any equilibrium that can be implemented via risk adjustment could also be implemented by setting the appropriate level of a uniform subsidy. Importantly, however, the converse is not true. That is, it is not the case that any equilibrium under a subsidy design can be implemented using appropriate risk adjustment. To see this, note that under risk adjustment, because insurers must set a single price, all buyers pay the same price, so equilibrium allocations must be monotone in willingness to pay. That is, if $v_i > v_k$ and individual k buys insurance in equilibrium then individual i also does. In contrast, non-uniform subsidies

[1] In many new markets, it is not uncommon to see such ex post adjustments that are based on realized costs. As markets mature, however, data availability allow for more accurate risk prediction and a more robust implementation of an ex ante risk adjustment system.

could be such that different consumer types face different prices, and if $s(w_k)$ is sufficiently greater than $s(w_i)$ the above monotonicity property could be violated, and the equilibrium may give rise to non-monotone insurance allocations.

2.3 Imperfect competition

We now consider a situation in which different buyers may have different valuations for different insurers. In such a situation, insurers have some amount of market power, and in a Bertrand-Nash Equilibrium markups are positive. Here we obtain the following result, where for simplicity we focus on the case of symmetric insurers, as in Mahoney and Weyl (2017).

Proposition 2 *Under imperfect competition and adverse selection, for any symmetric Nash equilibrium that is achievable with regular risk adjustment and no subsidies, in which markups are strictly positive, there is a (uniform) subsidy with no risk adjustment that leads to an equilibrium with the same enrollment for all types, and lower total spending for the market sponsor.*

The formal proof is in the appendix, but the intuition is the same as in Starc (2014) and Mahoney and Weyl (2017), and can be illustrated by examining the following insurer's first order condition:

$$p = ac(p) - \frac{q(p)}{q'(p)}\left(1 - ac'(p)\right), \tag{1}$$

where $q(\cdot)$ and $ac(\cdot)$ are the residual demand and residual average cost curves faced by the insurer. The key object in the first order condition is $ac'(\cdot)$, the derivative of the $ac(\cdot)$ curve with respect to the insurer's own price. With adverse selection ($ac'(\cdot) > 0$), the marginal buyer is cheaper than the average buyer and is therefore relatively attractive to cover, exerting downward pressure on prices and markups. Regular risk adjustments imply that the difference between the marginal and the average buyer is not as large (because the insurers are compensated more generously for covering high risk individuals), so $ac'(\cdot)$ is lower, thus reducing the pressure on prices and leading to greater markups. In other words, while in the perfect competition case the cost reductions implied by risk adjustment are fully passed through to consumers, when insurers have market power this pass through is incomplete, and subsidies are therefore a cheaper tool to lower consumer prices (and increase enrollment).

Proposition 2 implies that we can switch from any risk adjustment design to an environment with a uniform subsidy, leading to lower markups, greater coverage, and lower spending by the market sponsor. This has only considered uniform subsidies. In addition, the earlier observation described in last section remains. Risk adjustment (and/or uniform subsidies) imply monotone insurance allocations. Any non-uniform subsidy could relax this, and thus offers the market sponsor a greater set of equilibrium allocations that could be implemented.

2.4 Summary

Taken together, these results imply that subsidies lead to higher coverage (and consumer surplus) compared to risk adjustment for any given spending budget faced by the sponsor (or government). We call the enrollment difference between risk adjustment and uniform subsidies a "market power effect," by which subsidies lower insurers' markups for any coverage level that can be achieved. We call the additional enrollment gain from moving from uniform to non-uniform subsidies – and thus allowing the sponsor to provide different subsidies to different types of consumers – a "targeting effect," by which subsidies can be calibrated to provide more incentives for the low-risk types to enter the insurance pool.

The rest of the paper provides a quantitative assessment of these qualitative effects, using data and estimates from ACA Health Insurance Exchange in California. The quantitative analysis relaxes several of the simplifying assumptions of our stylized theory. In particular, we allow for insurers offering more than one plan, with vertically (as well as horizontally) differentiated products, and we allow for asymmetric costs across insurers.

3 Empirical setting

3.1 Setting and data

The health insurance exchange in California We use data from the first year (2014) of the California Health Insurance Exchange ("Covered California"), which was initiated by the Affordable Care Act (ACA). The Californian exchange is among the largest of the ACA exchanges, with over 1.5 million individuals obtaining health insurance coverage via this exchange during the first year. Given that over 90% of the buyers are subsidized and that the California exchange restricts substantially the scope of insurers' coverage design, it makes for a useful context in which we can quantitatively assess the interaction between risk adjustment and subsidies.

Covered California adopts all the regulations that are mandated by the national ACA reform, and also imposes several additional restrictions. We summarize below the key features; more detailed description is available in Tebaldi (2018). Covered California partitions the state into 19 geographic regions, each constitute a separate market. Insurers decide to participate in the market on a region-by-region basis. There are 3 to 6 insurers participating in each region. The two largest insurers, Anthem Blue Cross and Blue Shield of California, enroll 66% of individuals in Covered California, and this number grows to almost 90% when considering also HealthNet and Kaiser Permanente.

Within each region, participating insurers have to offer a set of four standardized coverage options. The standardized options are labeled by metals – bronze, silver, gold, and platinum – with each metal indicating a different level of coverage generosity, with approximate actuarial values that range from 60% (bronze) to 90% (platinum).[2] The details of each contract are also standardized,

[2] There is a fifth coverage level, "catastrophic coverage," which offers lower coverage than bronze. It is a high

as shown in Appendix Table A1. It is important to note that for Silver plans (and only for Silver plans), subsidized buyers receive cost-sharing subsidies in addition to premium subsidies. As shown in Appendix Table A1, for low income individuals these cost-sharing subsidies make silver coverage dominate the gold and sometimes even the platinum plan, even though the sliver premium is cheaper.

Despite these standardizations, insurers still differ along two important dimensions. They set different premiums as explained below, and they offer a different network of medical providers. The latter is an important source of unobserved heterogeneity across plans that our econometric model will account for. Fortunately, the provider network does not change with the level (metal) of the financial coverage.

Premium setting is subject to regulation, essentially constraining each plan-metal combination j in a given market (region) to set a single (base) price b_j. This base price is then mapped to the premium p_{ij} the individual has to pay and the subsidized portion of it s_{ij} that is paid by the government. These mappings are based on known, pre-specified (by the sponsor) formulas, which depend on the individual's age and income. Specifically, premiums are the product of the base price and an age factor $p_{ij} = f(age_i)b_j$, where $f(age)$ is a concave monotone function that is increasing from 1 (for 21 years old individuals) to 3 (for 64 years old individuals). The subsidy amount is then given by $s_{ij} = \max\{0, p_{ij} - \overline{p}_i\}$, where $\overline{p}_i = g(age_i, income_i)$ is set in a way that is benchmarked against the price the individual would have to pay for the second cheapest silver plan in the region.[3] Individual i thus pays $p_{ij} - s_{ij}$ out of pocket when selecting plan j. An important feature of this regulation is that buyers face identical choice sets but different prices, and because each plan sets a single base price the price variation across individuals is fully induced by the properties of the pre-specified formulas. This will prove useful for identification.

Data We use the same data as in Tebaldi (2018). We have the complete enrollment file of the California exchange for the 2014 coverage period; for each household enrolled in the exchange we observe the age and gender of its members, income, geographic location, premium paid, and plan selection. We combine these data with information on premiums, financial characteristics, and geographic availability for all the plans offered in the exchange. In addition, at the plan level (but not at the individual level) we observe the average (ex-post) realized amount of medical claims. Finally, in addition to these data on actual buyers, we use the American Community Survey (ACS) to construct a measure of potential buyers across different demographics in each region. We define a household as a potential buyer if it is either uninsured in 2013 or purchases insurance in the

deductible plan that is only available to individuals who are younger than 35 and who do not benefit of premium subsidies. These plans are not relevant for our analysis, so we abstract from them throughout the paper.

[3] This description applies for a single individual. Because income is typically measured at the household level, in practice the income and the subsidy formula are averaged over all covered individual within a household. The empirical exercise accounts for this additional complication, but we abstract from it in the text in order to ease notation and exposition.

individual (non-group) market in 2013.

Table 1 summarizes the characteristics of actual and potential buyers. Actual buyers tend to be relatively older and poorer. Several aspects shown in Table 1 are critical for our subsequent analysis. First, the majority of individuals (64%) select silver plans, presumably because – as mentioned earlier – cost-sharing subsidies make Silver plans substantially more attractive for subsidized households. Second, almost all actual buyers (94% of the households) are subsidized. Among subsidized buyers, 89% buy silver or bronze, which is why in our counterfactual analysis below we consider only silver and bronze plan enrollment.

Table 2 shows plan-level summary statistics. Two observations are important to highlight. First, with the exception of platinum plans, plan revenues are (on average) substantially above their incurred claims, but only due to the large amount of subsidies as enrollees' payments are substantially below plans' cost. Second, the data are consistent with the existence of asymmetric information, where higher coverage plans are associated with greater costs (Chiappori and Salanie 2000)[4]. This could reflect either adverse selection or moral hazard. A key focus of this paper will be the analysis of subsidies and risk adjustment in the presence of adverse selection; the empirical model we describe in the next section allows for such adverse selection, and the estimates provide direct evidence of adverse selection.

3.2 Empirical model

The baseline demand and cost estimates are the same as in Tebaldi (2018). For completeness we provide here a brief summary. A household i is associated with a vector of observable characteristics (age, income, region, and household size), denoted by z_i. A plan j is associated with its base price b_j, and a vector of observable characteristics (actuarial value, five insurer dummy variables, and region), denoted by x_j. For each household-plan combination we can then construct the implied premium p_{ij} and subsidy s_{ij}.

Demand We specify a mixed logit demand model, where the random utility coefficients are drawn from a finite support as in Berry, Carnall, and Spiller (2006) or Train (2008). Formally, the utility that household i derives from purchasing plan j is given by

$$u_{ij} = x_j' \beta_i - p_{ij} + \sigma_i \varepsilon_{ij}, \tag{2}$$

where $(\beta_i, \sigma_i) \sim_{iid} F(\beta, \sigma | z_i)$ and ε_{ij} is drawn, iid, from a type 1 extreme value distribution. The outside option is denoted by $j = 0$ and we adopt the standard normalization by which $x_0 = 0$ and $p_{i0} = 0$ for all i.

[4] This is true even after adjusting for the mechanical fact that for a given level of health care use, plans with greater coverage will report higher claims; since we know the difference in advertised actuarial value (see Appendix Table A1) this is a simple mechanical adjustment to make.

As mentioned, we assume that $F(\beta, \sigma|z_i)$ has a finite support that include four possible values $\{(\beta^k, \sigma^k)\}_{k=1}^4$, where $F(\cdot)$ follows a logistic distribution over this support:

$$F(\beta^k, \sigma^k|z_i) = \frac{\exp\left(\delta^k z_i\right)}{\sum_{l=1}^4 \exp\left(\delta^l z_i\right)}, \text{ for } k = 1, 2, 3, 4. \tag{3}$$

That is, the support of the random coefficient does not depend on the household characteristics, but the probability over this support does.

We allow demand to vary flexibly across markets (regions) by estimating it separately region by region. Overall, in each region there are 38 demand parameters: 28 parameters that define the support of $\left(\beta^k, \sigma^k\right)_{k=1}^4$, and 11 δ's (three δ's for each $k = 1, 2, 3, 4$, with one of these normalized to zero) that affect the distribution over this support as a function of household demographics.

Identification is discussed extensively in Tebaldi (2018). Loosely, the intuition for identification is that we only exploit the regulation-induced variation in prices of the same products within a region across buyers with different characteristics. We combine this with the functional form restrictions specified in (3), imposing constraints on how preferences can vary with age, income, and household size. In addition, identification of the value households place on the generosity of insurance (actuarial value) also relies on the discontinuity of cost-sharing subsidies at certain income thresholds (see Appendix Table A1, Panel B).

Cost We assume that the (expected) cost to the insurer of plan j from covering household i is given by

$$c_{ij} = \phi x_j + \gamma z_i + \omega_{ij}. \tag{4}$$

In this specification, ϕ captures cost differences across plans that are driven by differences in the geographic market, insurer, or generosity of insurance, holding the set of enrollees fixed. The parameter vector γ captures the extent to which insurers face different expected costs when enrolling households with different observable characteristics.

Finally, we assume that $\omega_{ij} = \lambda \nu_i + \eta_{ij}$, where η_{ij} is an iid error term while ν_i allows us to incorporate potential adverse selection by capturing the relationship between willingness to pay for higher coverage and plan cost. Specifically, we assume that ν_i is the household-specific coefficient on insurance generosity, which is a summary measure of the household's willingness to pay for greater coverage regardless of insurer choice. Adverse selection would be captured by $\lambda > 0$, which would imply that households with greater willingness to pay for coverage are associated with greater expected cost from coverage.

Because we observe choices at the household level but costs only at the plan level, in order to estimate equation (4) we aggregate it across all enrollees of a given plan. That is, for each j we define $\overline{z}_j$ and $\overline{v}_j$ as the average z_i and ν_i, respectively, across households enrolled in j. z_i is observed so aggregation is straightforward. In contrast, ν_i is unobserved, so we replace it with the posterior $E(v_i|i \text{ chose } j)$.[5] We then estimate

$$c_j = \phi x_j + \gamma \overline{z}_j + \lambda \overline{\nu}_j + \overline{\eta}_j \tag{5}$$

[5] That is, for each household we compute $E(v_i|i \text{ chose } j) = \sum_{k=1}^k \nu_i \frac{F(\beta^k, \sigma^k|z_i) \Pr[i \text{ chose } j|v_i]}{\Pr[i \text{ chose } j]}$.

via OLS.

3.3 Baseline estimates

Appendix Tables A2 and A3 provide details of the estimated demand and costs, respectively. These are discussed in detail in Tebaldi (2018). In Figure 2 we try to summarize the estimates by focusing on the key primitives that will enter our simulations of equilibrium under alternative designs of subsidies and risk adjustment schemes.

To construct the figure, we take every individual in the data and computes for her two summary measures of willingness to pay and expected cost. For willingness to pay, we compute the posterior estimate (conditional on actual choice) of the individual's parameter on actuarial value, and map it to the individual's willingness to pay for a 20 percentage point increase in the actuarial value of coverage (which is approximately moving individual from bronze to silver coverage for the average individual in the sample). For expected cost, we use equation (4) and assume that the individual is covered by Anthem's silver plan, which is the largest plan in the sample. The figure then reports the average and distribution of expected cost as a function of the individual's willingness to pay.

Figure 2 illustrates the extent of adverse selection: individuals with greater willingness to pay have higher expected costs. For example, individuals who are willing to pay approximately 500 dollars per-year to increase their coverage by 20 percentage points of actuarial value are (on average) expected to cost the insurer about 2,000 dollars. By contrast, individuals who are willing to pay as much as 2,000 dollars for the same increase in coverage have more than twice as high costs for the Anthem silver plan, approximately 4,600 dollars. The figure also illustrates the heterogeneity across individuals: as shown by the vertical range of the dashed lines in Figure 2, there is a fair amount of heterogeneity in the expected cost of individuals even conditional on their willingness to pay, which implies that selecting on demographics or other individual characteristics could have important implications for the nature and extent of adverse selection.

4 Quantitative estimates

4.1 Estimates for the California market

Setup We use the estimates of demand and cost in the California health insurance exchange to solve for market equilibria under alternative, counterfactual risk adjustment or subsidy schemes. For each buyer i, we generate an observable binary variable w_i that indicates a risk type. Specifically, in each region of Covered California, we compute the distribution of c_i and assume that whether the individual's cost is above or below the median (within the same region) is observable: $w_i = L$ if c_i is below the median, and $w_i = H$ otherwise. This is done for computational simplicity of our subsequent equilibrium simulations. Likewise, for simplicity, we assume that insurers offer only two contracts, bronze and silver. As noted, these are effectively the only two contracts purchased by subsidized enrollees, which is our focus. Because we allow for multiple (vertically differentiated) insurance contracts, the qualitative comparisons between risk adjustment and subsidies are not guaranteed by the theory, although in practice, as we will see, they appear to hold. Naturally the theory does not deliver quantitative estimates, which is the goal of this exercise.

To facilitate meaningful comparisons, in all equilibrium simulations we hold fixed across alternative market design regimes the total amount of spending by the market sponsor, which we denote by K. In what follows, we show results for $K = \$250$ per potential enrollee year and $K = \$500$ per potential enrollee year; this is roughly similar to the amount spent in Covered California.[6] For each value of K we assume that the objective is to maximize overall enrollment (alternatively, one could maximize consumer welfare, which leads to qualitatively similar results as shown in Appendix B) subject to the budget constraint and the specific restriction on the market design regime.

As described in Section 2, subsidies and risk adjustment consist of functions $s(w_i)$ and $r(w_i)$. With only two types, $w_i = L, H$, we can fully describe these policies by $s = (s^L, s^H)$ and $r = (r^L, r^H)$, which represent the subsidy and risk adjustment payments, respectively. As in Section 2, we restrict attention to regular risk adjustment, such that $r^H \geq r^L$.

We compute equilibria under four different market design regimes. The first solves for enrollment-maximizing risk-adjustment regime. That is, for each value of K, we set $s^L = s^H = 0$ and search for $r = (r^L, r^H)$ to maximize total enrollment subject to the budget constraint; we use the above estimates of demand and cost and assume a Bertrand-Nash equilibrium in premiums. The second regime considers a uniform subsidy. In this case we solve a similar problem, but set $s^L = s^H = s$ and $r^L = r^H = 0$, searching for the enrollment-maximizing value of s. Absent the ability to target subsidies (recall, $s^L = s^H$), the comparison between these two cases gives rise to the market power effect.

The two final regimes we consider allow for targeting and solve two different problems that allow subsidies (and hence premiums) to differ between L and H types. In one case, we impose a restriction that keeps all buyers (weakly) better off compared to the optimal flat subsidy. We thus solve the same optimization problem as in the flat subsidy regime, but allow for s^L and s^H to take different values and restrict overall enrollment of each type to be at lest as high as in the corresponding solution to the flat subsidy regime (with the same corresponding budget). This restrictions guarantees that (at the solution) all buyers are weakly better off than under the flat subsidy; we refer to this as the optimal Pareto subsidy. Finally, we move this last "Pareto" restriction, and solve for the optimal subsidies (s^L, s^H) that maximize overall enrollment without being concerned which type is being enrolled.

Results Table 3 shows the results. Consider first the results in Panel A ($K = \$250$). The first row shows results for the optimal risk adjustment policy. This policy pays insurers \$780 per enrollee-year for high risk enrollees (i.e. $r^L = 0; r^H = \$780$). It results in 43% enrollment, which does not

[6] In Covered California, this amount is approximately \$500 for the entire population and approximately \$300 for the population of subsidized single adults, which we focus on for our quantitative exercise. These figures are computed as the ratio between the total amount of subsidy outlays and the number of potential buyers as measured in the ACS.

vary much across types. The next row shows the gains from moving from optimal risk adjustment to an optimal uniform subsidy (i.e. $s^L = s^H = \bar{s}$). We estimate that this optimal uniform subsidy is \$480 per enrollee-year. The optimal uniform subsidy delivers a 4 percentage point (about 10 percent) increase in enrollment compared to optimal risk adjustment; low risk enrollment increases by about 3 percentage points, and high risk enrollment increases by about 4.5 percentage points.

The increase in enrollment with the uniform subsidy compared to risk adjustment is due to the markup effect discussed earlier. Risk adjustment increases markups because by definition it flattens the cost curve the insurer faces, making the marginal buyer (who, because of adverse selection, is healthier than the inframarginal buyer) less profitable relative to the inframarginal buyer than without risk adjustment. Indeed, as shown in Table 3, equilibrium markups under the uniform subsidy are \$161 (26 percent) lower than under risk adjustments. Coverage is higher given the lower markups, but overall producer surplus is almost 20 percent lower for precisely the pricing incentives highlighted in Proposition 2.

The third row shows the results from what we referred to as the optimal Pareto subsidy; we allow the subsidy to vary by type, while requiring that each types are weakly better off than under the flat subsidy. Relative to the optimal uniform subsidy, the optimal Pareto subsidy increases enrollment by about 1.5 percentage points (3 percent). The increase in enrollment from the optimal uniform subsidy is due to the targeting effect discussed earlier. In this example, the targeting effect is less than half the markup effect.

The optimal Pareto subsidy is \$660 per year for the low risk type, and \$420 per year for the high risk type. Relative to the optimal uniform subsidy, the subsidy rate for the low risk types has increased; this brings enough low-risk individuals into the market and sufficiently lowers the equilibrium premiums, allowing the market sponsor to lower the subsidy to the high risk type (thus respecting the budget constraint) and still have them (weakly) better off. Indeed, equilibrium premiums are lower, making all consumers better off. This is similar, conceptually, to Tebaldi (2018)'s analysis of optimal subsidies by age, in which lowering subsidies on the old and raising them on the young was able to make everyone (weakly) better off.

The last row of Panel A solves for the optimal subsidy without the Pareto restriction. In this example, interestingly, the result is identical to the optimal Pareto subsidy. In other words, the optimal subsidy in fact makes both types better off. This need not be the case, and indeed Panel B presents an example where it is not. Here, moving from the optimal Pareto subsidy to the optimal subsidy increases enrollment by another 0.4 percentage points, but does so at the cost of lowering enrollment for the high risk types. We suspect that this a function of the relative sizes of each cost type at different levels of willingness to pay for coverage. As a result, at lower levels of sponsor spending (hence higher premiums and lower enrollment), marginal enrollees are much cheaper than existing enrollees; the marginal enrollee therefore delivers a sufficiently large decrease in average costs that high risk individuals can be made better off despite a reduction in subsidy. As we go further down the distribution (i.e. through higher levels of sponsor spending and hence higher enrollment), the relationship between willingness to pay and costs flattens; as a result, the marginal enrollees have less of an impact on average costs; in order to enroll them, the higher risk individuals need to be made worse off.

The example illustrates the social tradeoff that can exist with optimal subsidies. Although optimal subsidies achieve higher enrollment (for the same level of sponsor spending) as risk adjustment, high risk types may be worse off under optimal subsidies than under optimal risk adjustment. Note that a move from optimal risk adjustment to the optimal Pareto subsidy involves no such tradeoff.

Figure 4 presents the main result for different levels of K under the four different market design regimes. The vertical distance between the different lines can be used as a metric to assess the different effects, as it represents the number of new enrollees we obtain "for free" (keeping the budget the same) as we move from a risk adjustment regime to a uniform subsidy regime, and then to a non-uniform (or targeted) subsidy regime. The first difference (between optimal risk adjustment and optimal uniform subsidy) corresponds to the market power effect, and the second (between optimal uniform subsidy and optimal non-uniform "Pareto" subsidy) to the targeting effect. As one can see, while the ranking of the regimes is, not surprisingly, consistent with the theory, the quantitative importance of each effect varies considerably with the overall enrollment share, which is closely related to the identity of the marginal enrollee.

The fact that the market power effect is increasing in spending (and enrollment) is again due to the shape of the cost curve in our setting. Recall (see Figure 2), that the correlation between willingness to pay and costs is higher for higher cost types. As a result, as enrollment increases, the slope of the average cost curve flattens and therefore the markup effect becomes larger. Put differently, at low levels of market spending (and hence enrollment), the enrollees are concentrated in the right hand side of Figure 2 (high willingness to pay) where the cost curve is much steeper; as a result, inframarginal enrollees are much less profitable than marginal enrollees which keeps markups lower.

4.2 A simplified setting: unpacking the black box

The above analyses provide a quantitative illustration of the difference in equilibria between optimal risk adjustment and optimal subsidies in the Covered California marketplace. In order to provide more intuition about the underlying economic forces behind these results, in this section we consider a substantially simplified setting, which corresponds to the simplifying assumptions from our stylized theoretical framework in Section 2.

Setup We consider the case of a duopoly, each offering only a single silver plan to single buyers with income between 138-400% of the Federal Poverty Level. We further assume that the two plans are symmetric, and only differ horizontally, with certain buyers preferring the network of physicians covered by A and others preferring the network of B. Formally, each potential buyer is associated with a two-dimensional vector of willingness to pay $v_i = (v_{iA}, v_{iB})$ and an identical (to both insurers) expected cost c_i, and the joint distribution of v_{iA} and v_{iB} and c_i is obtained from our baseline estimates (described in the last section), where (v_{iA}, c_i) and (v_{iB}, c_i) are drawn from our estimates for Anthem's silver plan, which is offered in every rating region in California. This translates to the aggregate and residual demand in the (simplified) market. The top left panel of Figure 3 plots the aggregate demand curve for insurance, as the prices of both insurers vary jointly ($p_A = p_B = p$, with p on the horizontal axis). The top right panel of Figure 3 depicts instead the residual demand curve faced by insurer B, holding p_A fixed at \$3.600.

This setting reflects a symmetric duopoly where firm have market power due to horizontal tastes but they do not systematically differ in terms of cost or attractiveness of their products. Combined with the distribution of v_i, our cost estimates imply the aggregate average and residual cost curve depicted in the top panels of Figure 3. Both aggregate and residual AC curves are downward sloping, which is a direct consequence of the presence of adverse selection in this market, as we already saw in Figure 2.

We continue to maintain our assumption of two observed types. The two bottom panels of Figure 3 illustrate how the aggregate and residual demand and cost curves, respectively, can be decomposed across low-risk and high-risk buyers. This highlights the primary source of adverse selection in this market: high-risk types are much more expensive to cover than low-risk types (on average, more than twice as much), and at the same time they have higher demand for insurance and lower sensitivity to premium changes. We continue to explore the same four market design regimes and to compare equilibria across them while holding fixed total market sponsor spending K.

Results In Section 2 we derived results that indicate that subsidies are "better" for two reasons. First, they lower markups since—compared to risk adjustment—they do not alter the relative profitability of marginal buyers; we called this a "market power effect." Second, they can limit inefficiencies arising from adverse selection by targeting different groups of buyers with different prices; we called this the "premium targeting effect." The simplified framework generates similar qualitative findings, but allows us to obtain clearer and graphical intuition as to the underlying theoretical objects that drive these overall effects.

Figure 5 allows us to illustrate the underlying "mechanics" behind the overall effect we discussed earlier. In the top panel, each point in the graph represents a pair of overall coverage for each type. The dashed, iso-WTP line represents the locus of coverage pairs that correspond to identical willingness to pay of the marginal buyer across both types. A risk adjustment or a flat subsidy regime forces equilibrium allocation to lie on this line, while a targeted subsidy regime frees up the market sponsor to find equilibrium allocation that are not on this line. The large black dots in Figure 5 represent the optimal allocation under the four different regimes for a budget of $K = 500$. Going form left to right, one can see that the market power effect allows more individuals of either type to get coverage because premiums (and markups) are lower under a flat subsidy relative to a risk adjustment regime. Yet, in both cases the allocation is forced to lie on the iso-WTP line. A targeted subsidy allows the market sponsor to offer a greater subsidy amount to low risk types, which in turn brings more of them to the market, reduces overall premiums, and can therefore can maintain similar equilibrium prices for high risk types (this specific channel is the focus of Tebaldi, 2018). Finally, this effect can be even larger in terms of overall enrollment if the market sponsor can "afford" some coverage reduction of high risk types, as shown in the right-most dot in Figure 5. The small gray dots in Figure 5 repeat a similar exercise for a lower budget (of $250). The qualitative nature of the exercise is similar, but the quantitative effect is different, with the market power effect and the relaxation of the Pareto constraint both being quite small.

The bottom-right panel of Figure 5 presents the corresponding subsidy amounts at each enrollment-maximizing level. the solid black line shows the optimal flat subsidy level, while the dashed-black lines and gray lines show the type-specific subsidy when targeting is allowed with and without the Pareto constraint (the higher subsidy is always associated with the low-risk types, who are the ones that are on the margin is most beneficial to attract into the market). The figure shows that at low levels of budget (and thus enrollment), the market is almost entirely comprised of high risk buyers, so attracting low risk buyers is important and optimal subsidies for them is very high. As coverage levels rise, the impact of attracting additional low-risk buyers into the market is lower, the value from targeting is smaller, and the subsidy amounts across types are not as different.

4.3 Extensions TBA (to be added in future drafts)

- We will show results with risk adjustment and subsidies combined
- We will show results with more "realistic" policies:
  - Use income, age, and actual risk score as observables w's
  - We will consider a case where subsidy is a function of age and income, while risk adjustment is a function of health risk score
- In the context of the simplified setting, we will explore comparative statics by showing results with
  - More granular risk scoring (richer than the current high/low binary indicator)
  - Less precise risk scoring (noisy signal about the current high/low indicator)

5 Conclusions

Our objective in the paper was to highlight that it makes sense to think jointly about subsidies and risk adjustment – two common market design instruments often employed by the same market sponsor – rather than to analyze each in isolation, as is typically done in both academic and health policy circles. Once we recognize that by shifting market demand and rotating market costs, respectively, subsidies and risk adjustment jointly interact to determine market equilibrium, the standard practice of thinking about subsidies as a way to achieve "affordability" and risk adjustment as a way to ameliorate adverse selection seems unsatisfactory.

We show theoretically that, at least under very stylized assumptions, subsidies can achieve greater enrollment for a given level of market sponsor spending. Using data and existing estimates from California's ACA health insurance exchange in 2014, we estimate that, holding sponsor spending fixed at roughly the level of current federal subsidy expenditures in this market, subsidies can increase enrollment by about 6 percentage points (13 percent) compared to optimal risk adjustment. Further, this increase in enrollment is achieved while holding all types (weakly) better off compared to the risk adjustment equilibrium.

A natural question is why, despite these theoretical and empirical results, risk adjustment remains an increasingly popular market design instrument. One possible economic explanation is that risk adjustment serves other functions beyond its role in pricing that we considered here. In particular, by decreasing the relative profitability to insurers of healthier compared to less healthy enrollees, risk adjustment may be important for reducing insurer cream-skimming efforts using non-price instruments, such as benefit design or marketing.[7] There are also potential political economy explanations. For example, while our theoretical and empirical analysis allows risk adjustment and subsidies to condition on the same type space, in practice offering greater subsidies to healthy consumers – as optimal subsidies often require – may conflict with naive intuition and may be politically difficult. Likewise, insurer profits may be higher under risk adjustment, creating a potential political force in favor of them. In this sense, our results can be thought of as providing a quantitative assessment of the costs of such potential constraints, in the context of California's ACA health insurance exchange.

More broadly, our intent here is not to prescribe specific market design strategies for health insurance exchanges, but rather to highlight the important sense in which two market design tools are highly related, and to provide some quantitative assessment of the tradeoff associated with greater reliance on risk adjustment relative to a richer and more flexible subsidy design.

References

Abraham, Jean, Coleman Drake, Dan Sacks, and Kosali Simon. 2017. "Demand for Health Insurance Marketplace Plans Was Highly Elastic in 2014-2015." .*Economics Letters* 159: 69-73.

Azevedo, Eduardo M., and Daniel Gottlieb. 2017. "Perfect Competition in Markets with Adverse Selection." *Econometrica* 85: 67–105.

Berry, Steven, Michael Carnall, and Pablo T. Spiller. 2006. "Airline Hubs: Costs, Markups and the Implications of Customer Heterogeneity." In *Advances in Airline Economics: Competition Policy and Antitrust*, Vol. 1, ed. Darin Lee, 183–214. Amsterdam: Elsevier.

Brown, Jason, Mark Duggan, Ilyana Kuziemko, and William Woolston. 2014. "How does Risk Selection Respond to Risk Adjustment? New Evidence from the Medicare Advantage Program." *American Economic Review* 104 (10): 3335-64.

[7] Another potential economic rationale for risk adjustment is that it allows conditioning payments on ex post realized costs, whereas subsidies must be based on ex ante measures. However, ex-post risk adjustment seems suboptimal, as it increases gaming opportunities, and this may be why we rarely see it in mature markets.

Chan, David, and Jonathan Gruber. 2010. "How Sensitive Are Low Income Families to Health Plan Prices?" *American Economic Review Papers and Proceedings* 100 (2): 292-296.

Curto, Vilsa, Liran Einav, Jonathan Levin, and Jay Bhattacharya. 2014. "Can Health Insurance Competition Work? Evidence from Medicare Advantage." NBER Working Paper No. 20818.

Decarolis, Francisco, Maria Polyakova, and Stephen P. Ryan. 2017. "The Welfare Effects of Supply-Side Regulations in Medicare Part D." NBER Working Paper No. 21298.

Dickstein, Michael J., Mark Duggan, Joe Orsini, and Pietro Tebaldi. 2015. "The Impact of Market Size and Composition on Health Insurance Premiums: Evidence from the First Year of the Affordable Care Act." *American Economic Review Papers and Proceedings* 105(2): 120–125.

Einav, Liran, Amy Finkelstein, and Mark R. Cullen. 2010. "Estimating Welfare in Insurance Markets using Variation in Prices." *Quarterly Journal of Economics* 125(3): 877–921.

Einav, Liran, Amy Finkelstein, Raymond Kluender, and Paul Schrimpf. 2016. "Beyond Statistics: The Economic Content of Risk Scores." *American Economic Journal: Applied Economics* 8(2): 195-224

Ellis, Randall P. 2008. "Risk Adjustment in Health Care Markets: Concepts and Applications." In M. Lu and E Jonsson (eds.), *Financing Health Care: New Ideas for a Changing Society*, Chapter 8. Wiley-VCH Verlag GmbH & Co., Weinheim, Germany.

Finkelstein, Amy, Nathaniel Hendren, and Mark Shepard. 2017. "Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts." Mimeo, MIT.

Frean, Molly, Jonathan Gruber, and Benjamin D. Sommers. 2017. "Premium Subsidies, the Mandate, and Medicaid Expansion: Coverage Effects of the Affordable Care Act." *Journal of Health Economics* 53: 72-86.

Geruso, Michael, and Timothy J. Layton. 2018. "Upcoding or Selection? Evidence from Medicare on Squishy Risk Adjustment." Mimeo, UT Austin.

Glazer, Jacob, and Thomas G. McGuire. 2000. "Optimal Risk Adjustment in Markets with Adverse Selection: An Application to Managed Care." *American Economic Review* 90 (4): 1055-71.

Handel, Ben, Igal Hendel, and Michael D. Whinston. 2015. "Equilibria in Health Exchanges: Adverse Selection versus Reclassification Risk." *Econometrica* 83: 1261–1313.

Jaffe, Sonia, and Mark Shepard. 2017. "Price-Linked Subsidies and Health Insurance Markups." Mimeo, Harvard University.

Timothy J. Layton, Thomas G. McGuire, and Anna D. Sinaiko. 2016. "Risk Corridors and Reinsurance in Health Insurance Marketplaces: Insurance for Insurers." *American Journal of Health Economics* 2(1): 66–95.

Mahoney, Neale, and E. Glen Weyl. 2017. "Imperfect Competition in Selection Markets." *Review of Economics and Statistics* 99(4): 637-651.

Newhouse, Joseph P., Mary Price, Jie Huang, J. Michael McWilliams, and John Hsu. 2012. "Steps To Reduce Favorable Risk Selection In Medicare Advantage Largely Succeeded, Boding Well For Health Insurance Exchanges." *Health Affairs* 31 (12): 2618-2628.

Starc, Amanda. 2014. "Insurer Pricing and Consumer Welfare: Evidence from Medigap." *RAND Journal of Economics* 45: 198–220.

Tebaldi, Pietro. 2018. "Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA." Mimeo, University of Chicago.

Train, Kenneth E. 2008. "EM Algorithms for Nonparametric Estimation of Mixing Distributions." *Journal of Choice Modelling* 1: 40–69.

Van de Ven, Wynand P.M.M., and Randall Ellis. 2000. "Risk Adjustment in Competitive Health Plan Markets." In A. Culyer and J. Newhouse (eds.), *Handbook of Health Economics* Vol. 1. Elsevier.

Appendix A: Proofs

In what follows, $q_j(p)$ denotes enrollment in insurer j's plan when prices are p, and $ac_j(p)$ the average cost of j when consumers face prices p. Additionally, we use $R_j(p)$ to indicate the per-enrollee risk adjustment transfer to insurer j when consumers face prices p. Adverse selection is defined as $\partial ac_j(p)/\partial p_j > 0$ for all j, so in the presence of adverse selection a regular risk adjustment implies $\partial R_j(p)/\partial p_j > 0$.

5.1 Proof of Proposition 1

Let $p^\star$ be an equilibrium prices when the risk adjustment $r(w)$ is adopted and $s(w) = 0$ for all w. Because $v_{i1} = v_{i2} = \ldots = v_{iJ}$ for all i, equilibrium is symmetric and insurers set the same price $p_1^\star = p_2^\star = ... = p_J^\star = \overline{p}$ and obtain the same risk risk adjustment transfer $R_j(p^\star) = \overline{R}$. Moreover, because insurers are identical, Bertrand-Nash equilibrium prices are set such that all insurers break even (as in Einav, Finkelstein, and Cullen, 2010):

$$\overline{p} = ac_j(p^\star) - \overline{R}. \tag{6}$$

Consider now the alternative policy in which there is no risk adjustment, while subsidies are

$$s(w) = \overline{s} = \overline{R}, \text{ for all } w. \tag{7}$$

The price $\widehat{p} = ac_j(p^\star)$ is then the new equilibrium, since $\widehat{p} = \overline{p} + \overline{s}$, and thus

$$\widehat{p} = \overline{p} + \overline{s} = ac_j(p^\star) - \overline{R} + \overline{s} = ac_j(\widehat{p} - \overline{s}), \tag{8}$$

so insurers break even. At this equilibrium, enrollment is the same for all types since net-of-subsidy prices are the same as in the original equilibrium, and the sponsor spending is the same since the per-enrollee payment in equation (7) is defined as the average risk adjustment payment under the original policy.□

5.2 Proof of Proposition 2

Consider first the Bertrand-Nash equilibrium with risk adjustment. Given the symmetric case we consider in the proposition, we focus on a symmetric equilibrium, such that $p_j^\star = p^\star$ and $R_j(p^\star) = R^\star$. In such an equilibrium, the following first order condition is satisfied:

$$p_j^\star = ac_j(p^\star) - R_j(p^\star) - \frac{q_j(p^\star)}{\partial q_j(p^\star)/\partial p_j}\left(1 - \frac{\partial ac_j(p^\star)}{\partial p_j} + \frac{\partial R_j(p^\star)}{\partial p_j}\right). \quad (9)$$

Consider now a case with no risk adjustment and a uniform subsidy that is given by

$$s^\star \equiv R_j(p^\star) + \frac{q_j(p^\star)}{\partial q_j(p^\star)/\partial p_j}\frac{\partial R_j(p^\star)}{\partial p_j}. \quad (10)$$

This level of subsidy is constructed so that it satisfies two key properties. First, it gives rise to a (symmetric) equilibrium in which each insurer j sets premium $\widehat{p}_j = p_j^\star + s^\star$. To see this, note that with subsidy $s^\star$ and no risk adjustment, equilibrium must satisfy the following first order condition

$$\widehat{p}_j = ac_j(\widehat{p} - s^\star) - \frac{q_j(\widehat{p} - s^\star)}{\partial q_j(\widehat{p} - s)/\partial p_j}\left(1 - \frac{\partial ac_j(\widehat{p} - s^\star)}{\partial p_j}\right). \quad (11)$$

Replacing $p_j^\star = \widehat{p}_j - s^\star$ implies

$$p_j^\star + s^\star = ac_j(p^\star) - \frac{q_j(p^\star)}{\partial q_j(p^\star)/\partial p_j}\left(1 - \frac{\partial ac_j(p^\star)}{\partial p_j}\right), \quad (12)$$

and substituting for $s^\star$ its construction from equation (10) yields the original first order condition from equation (9).

The second property of this particular construction of $s^\star$ is that

$$s^\star \equiv R_j(p^\star) + \frac{q_j(p^\star)}{\partial q_j(p^\star)/\partial p_j}\frac{\partial R_j(p^\star)}{\partial p_j}) < R_j(p^\star), \quad (13)$$

where the inequality follows from the fact that demand slopes down – $\frac{q_j(p^\star)}{\partial q_j(p^\star)/\partial p_j} < 0$ – and regular risk adjustments under adverse selection imply that $\frac{\partial R_j'(p^\star)}{\partial p_j} > 0$. This concludes the proof because it shows a subsidy design in which demand and insurance allocation remain the same as under risk adjustments, but sponsor expenditure is lower.□

Appendix B: Additional results

- Reproduce results when sponsor is maximizing consumer surplus rather than coverage

Figure 1: Intuition for Proposition 1

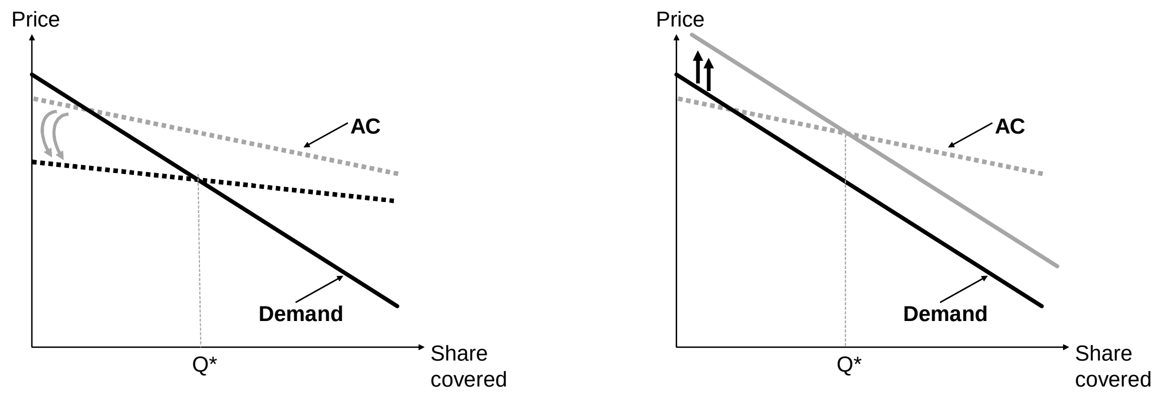

Figure provides intuition for the proof of Proposition 1. Figure plots demand and average cost curves, with competitive equilibrium given by the intersection of the two. In the left panel we illustrate the case of risk adjustment, which shifts and rotate the average cost curve. The right panel shows how a parallel shift in the demand curve, the result of a uniform subsidy, could achieve the same equilibrium allocation.

Figure 2: Summary of baseline estimates

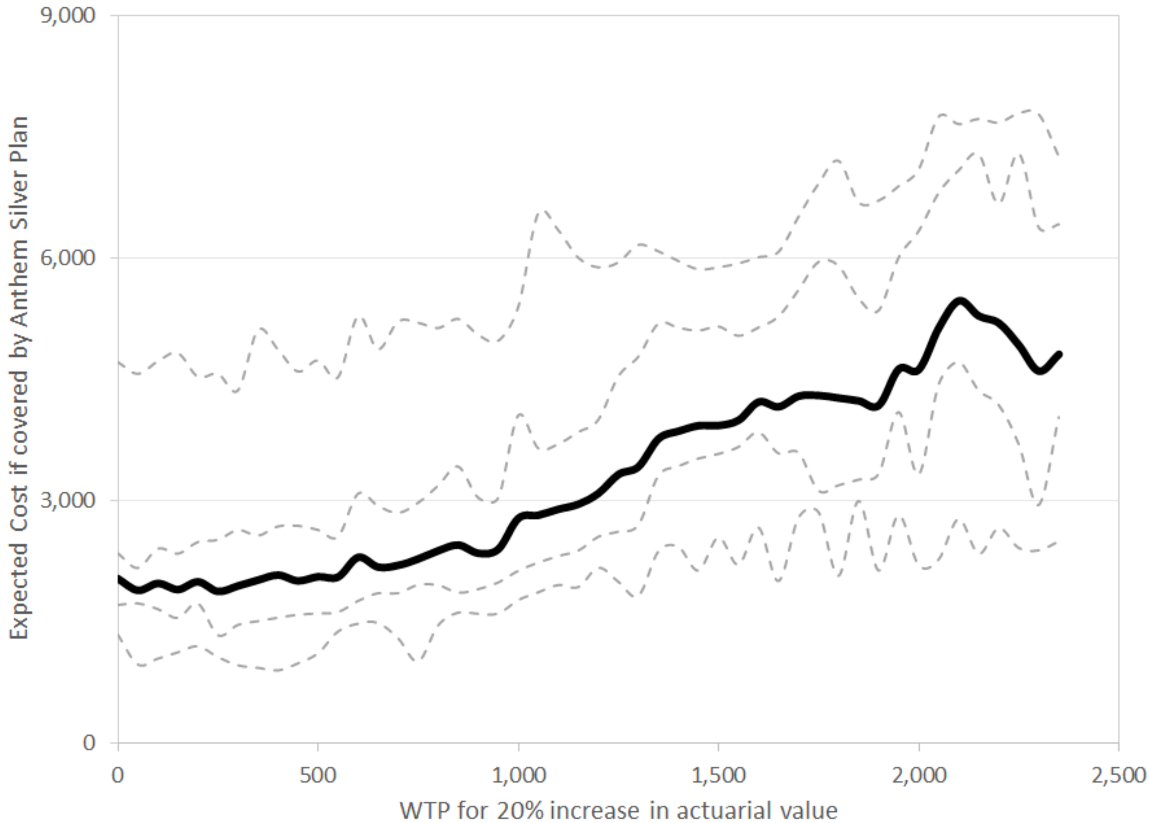

Figure uses the baseline estimates to plot the expected cost under Anthem silver plan (the largest plan) against the willingness to pay for a 20 percentage points increase in the actuarial value of the coverage. The solid black line present the average expected cost across all individual with a given willingness to pay, and the dashed gray lines present the 1st, 10th, 90th, and 99th percentiles in the distribution of individuals with a given willingness to pay. The upward sloping nature of the graph indicates the extent of adverse selection.

Figure 3: Demand and cost in the simplified setting

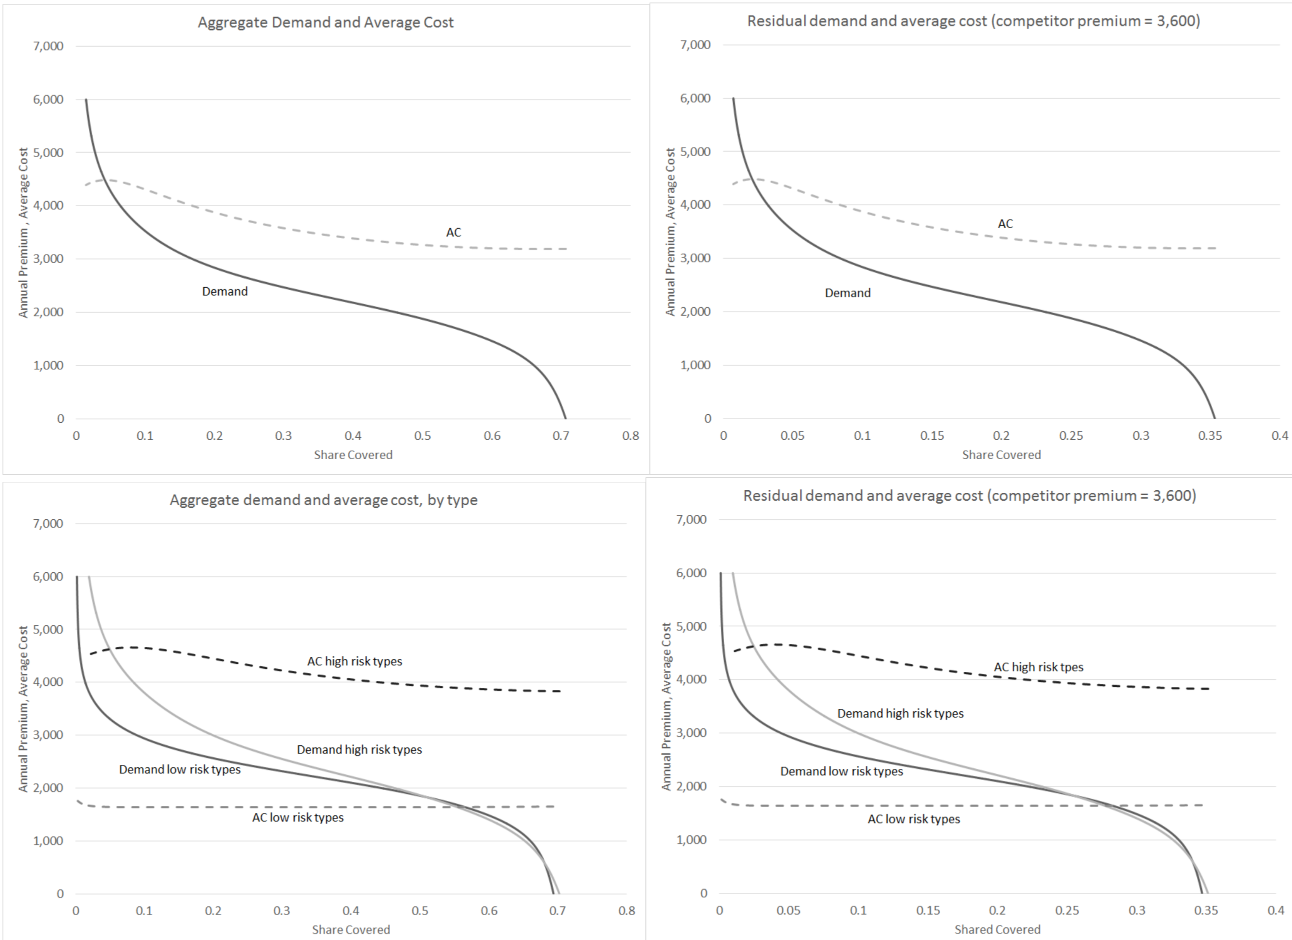

Figure illustrates the demand and average cost curves that are relevant for our simplified setting. In every figure, the horizontal axis indicates probability of purchase, while the vertical axis indicates annual premium or annual expected cost (in dollars). The top left panel plots aggregate demand and average cost if both insurers set the same premium, varying this premium between zero and \$6,000. In the bottom left panel the same curves are drawn separately for high-risk and low-risk consumers. The top right panel plots the residual demand and average cost curves faced by one insurer when the other insurer sets its annual premium at \$3,600. In the bottom right panel the same curves are drawn distinguishing between high-risk and low-risk consumers.

Figure 4: Total enrollment under different regimes

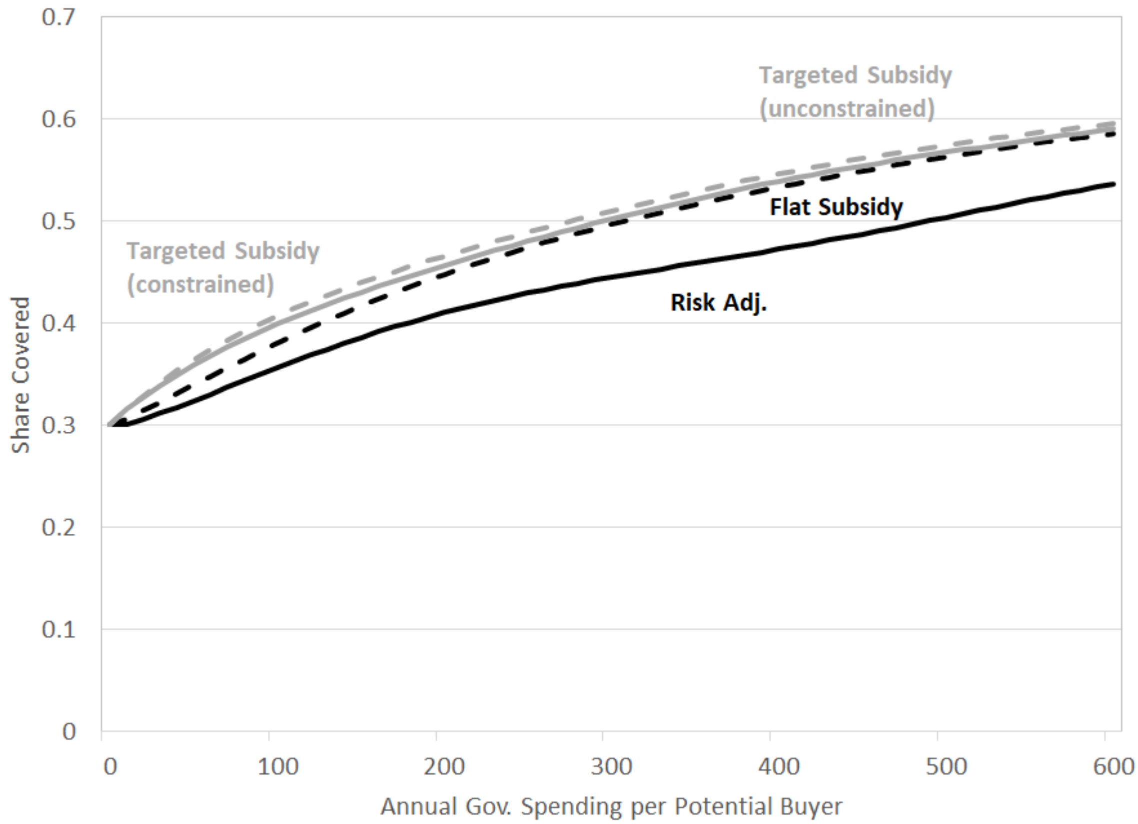

The Figure compares the total enrollment (expressed as share of potential buyers purchasing coverage on the vertical axis) as a function of the government spending (per potential buyer). Each line corresponds to a different regime, where for each regime we solve for the policy parameters that maximize total enrollment subject to the budget constraint. The vertical distance between the risk adjustment (solid black) line and the flat subsidy (dashed black) line indicates the "market power" effect, by which for any risk adjustment scheme there exists a flat subsidy scheme that achieves greater enrollment because insurers charge lower markups. The vertical distance between the flat subsidy (dashed black) line and the constrained targeted subsidy (solid gray) line indicates the "targeting" effect. This is the additional gain in terms of higher enrollment that subsidy can achieve, compared to risk adjustment, because the market sponsor can provide extra incentives for the participation of low-risk consumers. Finally, the dashed gray line corresponds to a situation in which targeted subsidies are optimized without requiring all buyers to be weakly better off than what they would be under the optimal risk adjustment policy

Figure 5: Unpacking the black box in a simplified setting

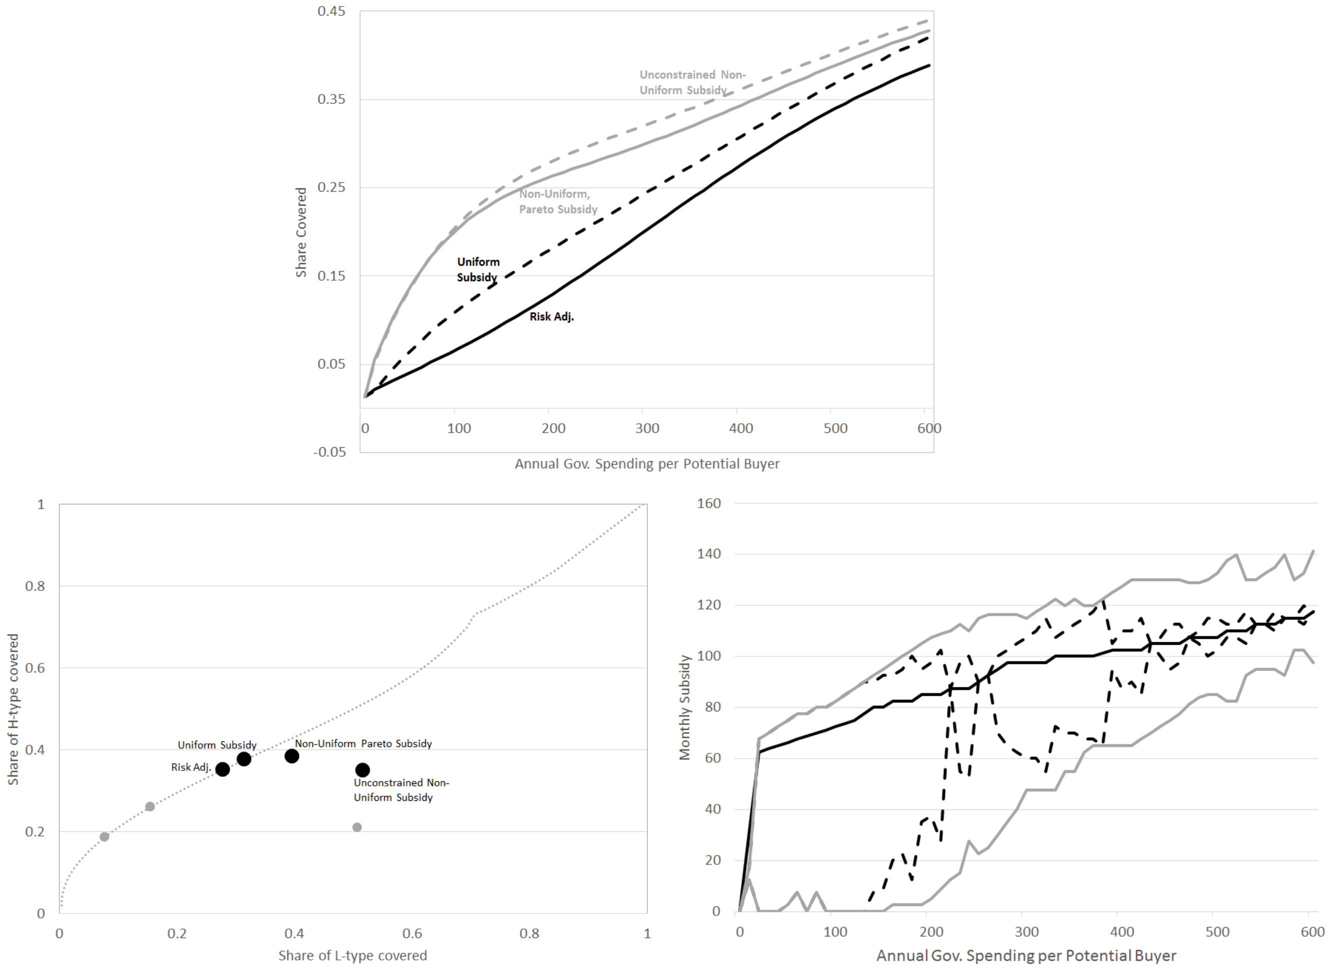

Top panel shows the allocations (in terms of enrollment) that corresponds to different regimes for two different budgets. The vertical and horizontal axes represent the share of potential buyers purchasing coverage among high-risk and low-risk individuals, respectively. The dotted line connects all possible allocations that can be achieved if both groups face the same premium. Any risk adjustment or flat subsidy policy implies allocations on this line. The labeled black dots indicate the allocations achieved by different regimes when the budget is \$500 per potential enrollee-year. The dots corresponding to the targeted subsidy are not on the dotted line because low-risk buyers face a lower net-of-subsidy premium. The smaller gray dots indicate the same regime combinations for a lower budget (of \$250 per potential enrollee-year). The bottom panel shows how the subsidies for the two types of consumers (vertical axis) vary as a function of the budget (horizontal axis). The solid black line shows the flat subsidy that maximizes enrollment. The dashed black lines are instead the optimal (constrained) target subsidies: the bottom line is the subsidy for high-risk buyers, and the top line is the subsidy for the low-risk buyers. Lastly, the solid gray lines correspond to the enrollment-maximizing targeted subsidies that allow high-risk buyers to be worse off than under the flat subsidy scheme. In this regime, the difference between the subsidies for the two types is larger.

Table 1: Summary statistics

	Observation	Mean
A. Potential buyers		
Age	3,392,942	41.34 (13.02)
Share subsidized (% FPL < 400)	3,392,942	0.66
Income as % of FPL (if subsidized)	2,231,013	231.65 (80.8)
Purchase coverage (households)	3,392,942	0.26
Purchase coverage (individuals)	6,122,167	0.21
B. Actual buyers		
Age	877,365	43.19 (12.98)
Share subsidized (% FPL < 400)	877,365	0.94
Income as % of FPL (if subsidized)	826,484	215.2 (63.2)
Bronze coverage	1,288,099	0.24
Silver coverage	1,288,099	0.64
Gold coverage	1,288,099	0.06
Platinum coverage	1,288,099	0.05

Table reports summary statistics for potential buyers and actual buyers in our data set; for continuous outcomes, standard deviations are reported in (parentheses). Panel A reports summary statistics for potential buyers using data from the 2013 ACS to describe the set of potential buyers in 2013. All estimates are reported at the household level except for when we report the share of individuals who purchase coverage. Panel B uses administrative data on enrollees in the California exchange in 2014. Demographics are reported at the household level (age is the average age of the household enrollees), while coverage is reported at the individual level (but shares are essentially identical at the household level).

Table 2: Plan pricing and enrollment

Tier	Observation	Mean	Std. Dev.	10th Pct	50th Pct	90th Pct
Panel A. Bronze						
Enrolees	82	3,764	3,575	434	2,692	8,726
Pre-subsidy premium	82	3,906	567	3,208	3,859	4,574
Post-subsidy premium	82	1,260	447	708	1,265	1,855
Average incurred claims	80	2,637	1,737	1,686	1,708	6,432
Panel B. Silver						
Enrolees	82	10,033	11,852	387	6,571	28,661
Pre-subsidy premium	82	5,303	860	4,181	5,307	6,500
Post-subsidy premium	82	1,522	585	873	1,421	2,325
Average incurred claims	80	3,497	1,284	2,562	2,948	6,432
Panel C. Gold						
Enrolees	82	940	1,148	61	618	2,503
Pre-subsidy premium	82	5,961	1,043	4,480	6,025	7,288
Post-subsidy premium	82	3,504	765	2,410	3,495	4,691
Average incurred claims	80	4,526	1,809	2,825	4,373	6,432
Panel D. Platinum						
Enrolees	82	822	916	33	573	1,829
Pre-subsidy premium	82	6,680	1,300	4,805	6,719	7,981
Post-subsidy premium	82	4,474	1,075	3,329	4,339	5,757
Average incurred claims	80	7,903	5,159	2,825	6,961	16,279

Table summarizes, for each coverage level, the number of enrollees, the per-person pre-subsidy premium received by the insurer, the per-person post-subsidy premium paid by buyers, and the per-person realized average cost. Each observation is an insurer-region pair, for a total of 82 plans. Two insurer-region observation are missing from the claims data: two small local insurers, Contra Costa and Valley, did not report claims for the 2014 coverage period.

Table 3: Outcomes under alternative market regimes

Policy	CS	PS	Markup	Risk Adjustment Low risk	Risk Adjustment High risk	Subsidy Low risk	Subsidy High risk	Share Covered Overall	Share Covered Low risk	Share Covered High risk	Share Silver Overall	Share Silver Low Risk	Share Silver High Risk
Panel A. Sponsor spending of $250 per potential enrollee-year													
Optimal Risk Adjustment	556.9	268.2	623.2	0	780	-------		0.430	0.403	0.436	0.285	0.400	0.652
Optimal Uniform Subsidy	608.1	217.9	461.9	-------		480	480	0.472	0.434	0.481	0.341	0.504	0.709
Optimal Non-Uniform Pareto Subsidy	638.1	197.5	404.2	-------		660	420	0.489	0.486	0.483	0.352	0.534	0.728
Optimal Unconstrained Non-Uniform Subsidy	638.1	197.5	404.2	-------		660	420	0.489	0.486	0.483	0.352	0.534	0.728
Panel B. Sponsor spending of $500 per potential enrollee-year													
Optimal Risk Adjustment	717.9	334.5	656.7	0	1,380	-------		0.509	0.476	0.516	0.337	0.372	0.652
Optimal Uniform Subsidy	810.7	254.8	456.5	-------		840	840	0.558	0.526	0.564	0.399	0.500	0.707
Optimal Non-Uniform Pareto Subsidy	836.5	255.6	448.7	-------		1,020	780	0.570	0.562	0.565	0.406	0.520	0.719
Optimal Unconstrained Non-Uniform Subsidy	837.3	251.2	438.4	-------		1,200	660	0.573	0.601	0.554	0.397	0.504	0.718

Table shows equilibria under different market design regimes. Panel A assumes a budget constraint by the market sponsor of $250 per potential enrollee-year; Panel B assumes a budget of $500 per potential enrollee-year. The first three columns show the consumer surplus (annual, per potential buyer), producer surplus (annual, per potential buyer), and the average markup (annual, per enrollee). The next four columns describe the optimal policy under each regime. The final columns report the share covered and share purchasing Silver (rather than Bronze), overall and for each consumer type.

Appendix Table A1: Coverage details

	Annual deductible	Annual max out-of-pocket	Primary visit	ER visit co-pay	Specialist visit co-pay	Preferred drugs co-pay	Advertised actuarial value[a]
Panel A. Insurance coverage before cost-sharing redusctions							
Bronze	$5,000	$6,250	$60	$300	$70	$50	60%
Silver	$2,250	$6,250	$45	$250	$65	$50	70%
Gold	$0	$6,250	$30	$250	$50	$50	79%
Platinum	$0	$4,000	$20	$150	$40	$15	90%
Panel B. Silver coverage after cost-sharing reductions							
Silver, >250% FPL	$2,250	$6,250	$45	$250	$65	$50	70%
Silver, 200-250% FPL	$1,850	$5,200	$40	$250	$50	$35	74%
Silver, 150-200% FPL	$550	$2,250	$15	$75	$20	$15	88%
Silver, 100-150% FPL	$0	$2,250	$3	$25	$5	$5	95%

Table describes the features associated with the different levels of coverage in the Covered California marketplace.

[a] Advertised actuarial values are computed by each insurer using a representative sample of claims provided by Covered California.

Appendix Table A2: Demand estimates

		Mean	10th Pct	25th Pct	Median	75th Pct	90th Pct
Panel A. Annual premium ($1,000s)							
Subsidized	no children, under-50	-1.580	-2.871	-2.314	-1.510	-0.849	-0.289
	no children, over-50	-1.482	-2.370	-1.802	-1.399	-1.062	-0.760
	with children	-0.442	-1.064	-0.751	-0.341	-0.046	0.077
Unsubsidized		-1.553	-2.827	-2.218	-1.529	-0.797	-0.232
Panel B. Actuarial value (%)							
Subsidized	no children, under-50	0.078	0.011	0.036	0.082	0.118	0.141
	no children, over-50	0.113	0.051	0.073	0.115	0.147	0.169
	with children	0.046	0.005	0.014	0.036	0.071	0.103
Unsubsidized		0.079	0.021	0.044	0.084	0.115	0.130
Panel C. Difference between Blue Cross Blue Shield and Anthem							
Subsidized	no children, under-50	0.242	-1.127	-0.374	0.452	1.140	1.720
	no children, over-50	0.240	-1.375	-0.449	0.607	1.278	1.732
	with children	0.541	-0.765	-0.276	0.953	1.601	1.877
Unsubsidized		0.202	-0.680	-0.330	0.386	1.067	1.611
Panel D. Difference between HealthNet and Anthem							
Subsidized	no children, under-50	-0.349	-2.324	-1.250	0.049	0.416	0.665
	no children, over-50	-0.453	-1.680	-1.143	-0.157	0.133	0.291
	with children	0.106	-1.541	-0.233	0.373	0.735	1.083
Unsubsidized		-0.446	-2.096	-1.183	-0.156	0.223	0.731
Panel E. Difference between Kaiser and Anthem							
Subsidized	no children, under-50	-0.108	-1.727	-0.363	0.346	0.717	1.154
	no children, over-50	0.042	-1.854	0.058	0.544	0.987	1.370
	with children	-0.161	-1.847	-0.574	0.323	0.750	1.015
Unsubsidized		-0.023	-1.736	-0.140	0.340	0.744	1.103
Panel F. Difference between other minor insurers and Anthem							
Subsidized	no children, under-50	-1.370	-2.141	-1.973	-1.721	-1.096	-0.060
	no children, over-50	-1.708	-2.387	-2.294	-2.102	-1.185	-0.455
	with children	-1.365	-2.335	-2.026	-1.656	-1.219	0.691
Unsubsidized		-1.386	-2.217	-1.994	-1.742	-0.668	-0.051

Table shows summary statistics of random coefficients based on mixed logit estimates. For each parameter and demographic group, the table shows the average of the corresponding coefficient, as well as 10th, 25th, 50th, 75th, and 90th percentiles. Within each group, the parameters vary across regions, and across different combinations of age, income, and household size.

Appendix Table A3: Cost estimates

	Log-annual claims						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A. Product characteristics							
Bronze	(Omitted)	(Omitted)	(Omitted)	(Omitted)	(Omitted)	(Omitted)	(Omitted)
Silver	0.478***	0.399***	0.344**	0.432**	0.370***	0.306**	0.411**
	(0.132)	(0.144)	(0.140)	(0.170)	(0.141)	(0.140)	(0.170)
Gold	0.648***	0.619***	0.597***	0.608***	0.604***	0.582***	0.598***
	(0.078)	(0.079)	(0.077)	(0.085)	(0.080)	(0.075)	(0.083)
Platinum	1.219***	1.171***	1.158***	1.189***	1.130***	1.113***	1.153***
	(0.098)	(0.100)	(0.097)	(0.107)	(0.104)	(0.097)	(0.106)
Panel B. Buyer characteristics							
Age	0.0188**	0.0187**	0.0206**	0.0221**	0.0187**	0.0198**	0.0217**
	(0.00846)	(0.00838)	(0.00830)	(0.00911)	(0.00803)	(0.00817)	(0.00889)
FPL	0.00277	0.00243	0.00207	0.00321	0.00275	0.00227	0.00351
	(0.00251)	(0.00253)	(0.00236)	(0.00297)	(0.00245)	(0.00234)	(0.00298)
Household	0.087	0.057	0.144	0.199	0.057	0.138	0.185
	(0.112)	(0.111)	(0.131)	(0.153)	(0.115)	(0.135)	(0.157)
Low WTP (<1700)		(Omitted)	(Omitted)	(Omitted)			
High WTP (>1700)		0.113	0.172**	0.123			
		(0.078)	(0.077)	(0.089)			
WTP in [0,1100]					(Omitted)	(Omitted)	(Omitted)
WTP in (1100,1400]					0.078	0.096	0.094
					(0.105)	(0.096)	(0.116)
WTP in (1400,1700]					0.133	0.124	0.099
					(0.109)	(0.100)	(0.118)
WTP in (1700,2000]					0.260**	0.291***	0.243*
					(0.115)	(0.109)	(0.132)
WTP > 2000					0.147	0.263**	0.173
					(0.126)	(0.121)	(0.137)
Region FE	Y	Y	Y	Y	Y		
Insurer FE	N	N	Y	N	Y		
Insurer-Regior	N	N	N	Y	N	N	Y
Observations	338	338	338	338	338	338	338
R-squared	0.463	0.467	0.572	0.597	0.472	0.575	0.600

Table shows parameters of the cost model estimated from equation (4). Each observation is an insurer-region-tier triplet, where I exclude Catastrophic coverage since it is not available for subsidized enrollees. After this exclusion, claims data used for estimation cover over 90% of enrollment, with two missing carriers (Contra Costa and Valley). Buyer characteristics are computed as average across enrollees of the plan, where WTP is the posterior of the ratio β/α , conditional on observed choice, based on mixed logit estimates. Robust standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1